



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

Address: COMMISSIONER FOR PATENTS

P.O. Box 1450

Alexandria, Virginia 22313-1450

www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/717,730	11/20/2003	Daniel N. Cripe	200313587-1	8712
22879 7590 12/24/2009 HEWLETT-PACKARD COMPANY Intellectual Property Administration 3404 E. Harmony Road Mail Stop 35 FORT COLLINS, CO 80528				
EXAMINER SINKANTARAKORN, PAWARIS				
ART UNIT 2464		PAPER NUMBER		
NOTIFICATION DATE 12/24/2009		DELIVERY MODE ELECTRONIC		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM

ipa.mail@hp.com

laura.m.clark@hp.com



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/717,730
Filing Date: November 20, 2003
Appellant(s): CRIPE ET AL.

Nick P. Patel
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed *September 29, 2009* appealing from the Office action mailed *July 27, 2009*.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

USPN 6,151,297	Congdon	11-2000
USPN 6,938,092	Burns	8-2005
USPN 6,314,525	Mahalingham	11-2001

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 103

This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-4 and 6-21 are rejected under 35 U.S.C. 103(a) as being unpatentable over Congdon et al. (USPN 6,151,297) in view of Burns (USPN 6,938,092).

Regarding claim 1, Congdon et al. disclose a computer system comprising:
a central processing unit (CPU) (see column 6 line 15, Network Operating System runs on a CPU);

a first and second network adapter teamed together and configured to receive offloaded connections (see column 8 lines 57-64, when multiple NICs in a server are attached to a network and the NICs are using the same MAC address, it is possible to receive packets on many ports) ; and

wherein a program executing on the CPU reloads an offloaded connection established by the first network adapter onto the second network adapter if one of a plurality of packets associated with the offloaded connection was received on the second network adapter (see column 8 lines 1-5, 12-14, and 26-39, the NICs are active on the network at the same time and the invention supports a fault tolerance feature; fault tolerance enables a system to continue operating properly in the event of the failure of some of its components. The switch selects one of the multiple NICs using the fault tolerance feature when one of the NICs fails so that there is no need to reestablish a new connection).

Congdon et al. merely disclose fault tolerance. However, Burns, from the same or similar fields of endeavor, discloses a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter (see column 6 line 36 – column 7 line 29 and column 7 lines 57-67, Port 2 becomes a primary port, which is responsible for receiving data

when Port 1 fails; as a result of Port 2 being a primary port, which is being used to receive data, the port aggregation driver updates the TCB for TCP connection #1 by changing MAC C to MAC A, where the port aggregation driver is of network access layer (see Figure 2 reference numeral 115); Once the MAC addresses are swapped, the packets originally destined to port 1 are received at port 2, which now has MAC A as its address; as a result of the swap of the MAC addresses, port aggregation driver calls a NDIS request function to update the handle and pointer of the change and hence reloads an offloaded connection).

Thus, it would have been obvious to the person of ordinary skill in the art at the time of the invention to implement a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter of Burns into the switch of Congdon et al.

The motivation for implementing a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter is that it increases the efficiency of the switch by implementing fault recovery feature into the switch.

Regarding claims 2 and 17, the first and second network adapters are capable of fully offloading all protocol processing (see column 8 lines 9-14, if one of the NICs fails, there is no need to reestablish a new connection; therefore, the NICs are fully capable of offloading all protocol processing);

regarding claims 3 and 18, the first and second network adapters transmit and receive packets of data using a single media access control (MAC) and internet protocol (IP) address (see column 8 lines 1-5).

Regarding claim 4, Congdon et al. fail to teach the system, wherein the program reloads an offloaded connection by transferring the context of the connection from the first network adapter to the second network adapter. However, Burns from the same or similar fields of endeavor discloses a method of the program reloads an offloaded connection by transferring the context of the connection from the first network adapter to the second network adapter (see column 6 lines 36-67 and column 7 lines 57-67, the port aggregation driver updates the TCB for TCP connection #1 by changing MAC C to MAC A).

Thus, it would have been obvious to the person of ordinary skill in the art at the time of the invention to implement a method of the program reloads an offloaded connection by transferring the context of the connection from the first network adapter to the second network adapter of Burns into the switch of Congdon et al.

The motivation for implementing a method of the program reloads an offloaded connection by transferring the context of the connection from the first network adapter to the second network adapter is that it increases the efficiency of the switch.

Regarding claims 6 and 21, Congdon et al. fail to teach the system, wherein the first and second network adapters send a notification to the program if a connection is prematurely terminated. However, Burns from the same or similar fields of endeavor discloses a method for send a notification to the program if a connection is prematurely

terminated (see column 6 lines 36-45, sending flush status to the port aggregation driver).

Thus, it would have been obvious to the person of ordinary skill in the art at the time of the invention to implement a method for send a notification to the program if a connection is prematurely terminated of Burns into the switch of Congdon et al.

The motivation for implementing a method for send a notification to the program if a connection is prematurely terminated is that it increases the efficiency of the switch.

Regarding claim 7, Congdon et al. disclose a system, wherein the first and second network adapters comprise network interface cards (NICs) (see column 6 line 2).

Regarding claims 8, 12, and 16, Congdon et al. disclose a method comprising:
examining a packet received from an external device 9(see column 7 lines 38-42);
determining whether a connection associated with the packet is currently offloaded (see column 8 lines 26-39, fault tolerance feature determines whether there is any fail NIC in the server, therefore, determines whether a connection associated with the packet is currently transmitted);

reloading the connection if the packet associated with the connection is offloaded and received by a network interface not currently processing the offloaded connection (see column 8 lines 1-5, 12-14, and 26-39, the NICs are active on the network at the same time and the invention supports a fault tolerance feature; fault tolerance enables a system to continue operating properly in the event of the failure of some of its components. The switch selects one of the multiple NICs using the fault tolerance

feature when one of the NICs fails so that there is no need to reestablish a new connection).

Congdon et al. merely disclose fault tolerance. However, Burns from the same or similar fields of endeavor discloses a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter (see column 6 line 36 – column 7 line 29 and column 7 lines 57-67, Port 2 becomes a primary port, which is responsible for receiving data when Port 1 fails; as a result of Port 2 being a primary port, which is being used to receive data, the port aggregation driver updates the TCB for TCP connection #1 by changing MAC C to MAC A, where the port aggregation driver is of network access layer (see Figure 2 reference numeral 115); Once the MAC addresses are swapped, the packets originally destined to port 1 are received at port 2, which now has MAC A as its address; as a result of the swap of the MAC addresses, port aggregation driver calls a NDIS request function to update the handle and pointer of the change and hence reloads an offloaded connection).

Thus, it would have been obvious to the person of ordinary skill in the art at the time of the invention to implement a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter of Burns into the switch of Congdon et al.

The motivation for implementing a method of reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter is that it increases the efficiency of the switch by implementing fault recovery feature into the switch.

Regarding claims 9 and 13, Congdon et al. disclose a method further comprising determining an identifier for the network interface that receives the packet (see column 7 lines 38-44, the switch determines the output port by looking up the Destination Address in the address table) and writing the determined identifier to a memory (see column 7 address table).

Regarding claims 10, 14, and 19, Congdon et al. fail to teach the method, wherein the reloading further comprising copying the context of the connection to the network interface that received the packet. However, Burns from the same or similar fields of endeavor discloses the method, wherein the reloading further comprising copying the context of the connection to the network interface that received the packet (see column 6 lines 36-67 and column 7 lines 57-67, the port aggregation driver updates the TCB for TCP connection #1 by changing MAC C to MAC A).

Thus, it would have been obvious to the person of ordinary skill in the art at the time of the invention to implement the method, wherein the reloading further comprising copying the context of the connection to the network interface that received the packet of Burns into the switch of Congdon et al.

The motivation for implementing the method, wherein the reloading further comprising copying the context of the connection to the network interface that received the packet is that it increases the efficiency of the switch.

Regarding claims 11 and 15, Congdon et al. disclose a method, wherein the network interface that received the packet and the network interface currently offloading the connection are teamed together (see column 6 lines 13-17, the group of NICs appear as a single NIC to the clients in the network);

regarding claim 20, the program monitors all data received by the first and second means for sending and receiving data connections (see column 7 lines 38-44, the switch determines the output port by looking up the Destination Address in the address table).

Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over Congdon et al. in view of Burns as applied to claims and 1 and 16 above, and further in view of Mahalingham et al. (USPN 6,314,525).

Regarding claim 5, Congdon et al. in view of Siu et al. disclose all the subject matter of the claimed invention except the system/method, wherein the program inactivates connections associated with packets that have not been received for a defined time period.

However, the invention of Mahalingham et al. from the same or similar fields of endeavor disclose a method/system for deactivating a network adapter when the

network adapter fails to respond after a predetermined time period (see column 9 lines 46-56).

Thus, it would have been obvious to the person of ordinary skill in the art to implement a method/system for deactivating a network adapter when the network adapter fails to respond after a predetermined time period as taught by Mahalingham into the data processing method of Congdon et al.

The motivation for implementing the method/system for deactivating a network adapter when the network adapter fails to respond after a predetermined time period is that it increases efficiency of the NICs in the server.

(10) Response to Argument

Regarding claims 1-4 and 6-21, The Appellants argue that Burns does not teach that port 2 is selected as a result of receiving a packet associated with the connection previously coupled to port 1; thus, Burns does not teach reloading an offloaded connection established by the first network adapter onto the second network adapter **as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter.**

In response to Appellants' argument, The Examiner respectfully disagrees with the argument above.

Burns teaches, in normal configuration, a TCP offload connection is handled in Fast-Path by a TCP offload network interface device (NID) between multiple ports of an aggregation team (see abstract), where port 1 having MAC address A is used for

receive and port 2 having MAC address B is used for transmit (see Figure 3 and Figure 4 step 400).

However, when failure occurs in port 1, port aggregation driver causes the network interface device to change the MAC address of port 1 from MAC address A to MAC address C, and to change the MAC address of port 2 from MAC address C to MAC address A (see Figures 3 and 4 and column 6 line 36 – column 7 line 13). Thus, port 2 is now used for receive as well as for transmit, where packets with destination address of MAC address A will be diverted to port 2, rather than port 1.

Subsequently, following port 2 being used in both receive mode and transmit mode, if transmit for the TCP connection is to remain on port 2 and if port 2 is now associated with MAC address A, then the MAC source address in Transmit Control Block (TCB) must be changed from MAC address C to MAC address A. Accordingly, port aggregation driver updates the TCB for TCP connection by changing the MAC source address to be the MAC address of the selected member chosen for transmission, which is port 2 in this particular case (see Figures 3 and 4 step 406 and column 7 lines 57-67).

Thus, the original cause of the updating of the TCB for TCP connection occurs is the failure of port 1, where the failure of port 1 results in port 2 receiving packets associated with the offloaded connection of port 1 (see Figure 4, step 401, step 403, and step 406). And, as a result of the failure of port 1 and port 2 receiving packets associated with the offloaded connection of port 1, Burns teaches, in the flow chart in Figure 4, the step 406 of updating the TCB for TCP connection by changing the MAC

source address to be the MAC address of port 2 occurs. Step 406 occurs as a result of series of steps 401 through steps 405 occurring. On the other hand, Step 406 would not have occurred if steps 401 through steps 405 did not occur. Thus, Burns teaches reloading an offloaded connection established by the first network adapter onto the second network adapter as a result of one of a plurality of packets associated with the offloaded connection being received on the second network adapter.

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

Conclusion

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Pao Sinkantarakorn/

Examiner, Art Unit 2464

Conferees:

/Ricky Ngo/

Supervisory Patent Examiner, Art Unit 2464

/Huy D Vu/

Supervisory Patent Examiner, Art Unit 2461